

# Are we there yet? Investigating solution convergence in a high-resolution global ocean model with a 2048 processor cluster of distributed shared memory SGI Altix machines.

CHRIS HILL, DIMITRIS MENEMENLIS, BOB CIOTTI, CHRIS HENZE

General Terms: Computational fluid dynamics, ocean modeling, parallel computing

---

## 1. INTRODUCTION

In this article we describe technical aspects of global ocean model configurations with resolutions up to 5km ( $\approx \frac{1}{16}^\circ$ ) that exploit a testbed 2048 Itanium-2 processor SGI Altix system at the NASA Ames Research Center. The size of this system makes possible global ocean simulations at resolutions that until now have been impractical. The calculations we describe employ an Adams-Bashforth time-stepping procedure on a finite volume grid with up to 1.25 billion grid cells. This workload is spread evenly over 1920 of the Altix processors so that each individual processor is responsible for simulating around 586,000 grid cells (corresponding to a surface region roughly 210 x 210 km in extent). Decomposing the workload over this many processors yields a setup that, with extensive diagnostics and analysis options included, uses about 870MB of main memory per processor and can integrate forward at a rate of around  $5\mu\text{s}/\text{timestep}/\text{gridcell}$ . With a timestep of two minutes this performance allows a year of simulation to be completed in under ten days.

The model configurations we employ are significant in that, at the resolutions the Altix system makes possible, numerical ocean simulations begin to truly represent the key dynamical process of oceanic meso-scale turbulence. Meso-scale turbulence in the ocean is, in some ways, the analog of synoptic weather fronts in the atmosphere. However, because of the density characteristics of seawater, the length scale of turbulent eddy phenomena in the ocean is around ten or less kilometers. In contrast, in the atmosphere, where the same dynamical process occurs, it has length scales of thousands of kilometers - see figure (1). Although it has been possible to resolve ocean eddy processes well in regional ocean simulations [Hurlburt and Hogan(2000)] and process studies [Jones and Marshall(1997); Haine and Marshall(1998); Nurser and J.W.Zhang(2000)] for some time, global scale simulations that resolve or partially resolve the ocean's energetic eddy field are still rare [Maltrud and McClean(2005); Masumoto et al.(2004)] because of the immense computational challenge they represent.

Physically ocean meso-scale eddies contribute to regulating large-scale ocean transports of momentum, heat and fresh water, all of which play an important role in creating the Earth's climate. Eddies also interact, non-linearly, with the larger scale general circulation thereby modifying an ocean simulation's mean state




Fig. 1. Atmosphere and ocean eddies. Eddies occur in both the atmosphere and ocean but with vastly different length scales. The left panel shows a polar view of the atmospheric jet stream meandering as it circulates the northern hemisphere. The length scales of the meanders are several thousand kilometers. The right panel shows satellite imagery of Gulf Stream meanders off the U.S. East Coast. In places the length scales of the meanders are as small as ten kilometers.

and altering its transient response to perturbations. Through these processes, and through the processes by which ocean energy is dissipated at basin margins, meso-scale eddies play a significant, complex, and non-linear role in climate. Process studies [Jones and Marshall(1997)] suggest that resolutions of a few kilometer are needed to adequately resolve eddy processes. Consequently ocean modelers are continually on the look out for technological means to increase the resolution of global-scale integrations in order to converge on an adequate representation of these non-linear behaviors. The work we describe here creates the technological machinery for examining how much resolution is needed for eddies, and their impacts, to be accurately represented in a global ocean model.

The article is organized as follows. We first review prior work in section 2. A description of the underlying mathematical and numerical basis of our model together with a discussion of its computational formulation is in section 3. Section 4 provides details on how the 2048 shared memory Altix cluster is configured, looks at how our parallel algorithm maps to the system, and describes the resultant performance. In sections 5 and 6, we summarize the prospects for using a large Altix system for optimal representation of eddy terms in a global ocean model.

## 2. PRIOR WORK

Eddying calculations first started to appear around 1990 [Bryan and Holland(1989); Boening and Budich(1992)]. These initial calculations were facilitated by the availability of vector machines such as the Cray YMP with sustained memory bandwidths of several gigabytes/second per processor. Global eddying calculations were first carried out [Semtner and Chervin(1992)] around this time too. During this period computational resources supported simulations that permitted eddies, but did not allow models to resolve the scales shown in figure 1. More recently

[Hurlburt and Hogan(2000)], [Smith et al.(2000)], [Oschillies(2002)] and [McClean et al.(2002)] have undertaken large-scale regional simulations that are termed eddy-resolving in that they resolve the first baroclinic Rossby radius of deformation in the ocean [Chelton et al.(1998)], a physical estimate of a typical ocean eddy size. In these regional studies authors find that the transition from eddy-permitting to eddy-resolving produces qualitative improvements in the fit between simulations and observations. Following these and other studies [Maltrud and McClean(2005)] reported on a global  $\frac{1}{10}^\circ$  simulation on an IBM Power 3 system and [Masumoto et al.(2004)] described a global  $\frac{1}{10}^\circ$  simulation on the Earth Simulator in Japan. These studies further reinforce the notion that resolving eddy processes provides a significant boost in model skill. In this paper we report on the computational aspects of a study that uses a cluster of four 512-way NUMA Altix systems for global eddy resolving simulations at up to  $\frac{1}{16}^\circ$  resolution. The study is aimed at developing a clearer understanding of processes underlying skill improvements that eddy resolving models show. To develop this understanding we perform a series of heavily instrumented simulations that produce more than a terabyte of output for each year simulated and that repeatedly process more than one billion grid cells in a time-stepping procedure. In doing this we gain insights into the performance of a clustered Altix system with 2048 CPUs for ocean simulation.

### 3. ALGORITHM

The Massachusetts Institute of Technology General Circulation Model (MITgcm) algorithm is rooted in the incompressible form of the Navier-Stokes equations for fluid motion in a rotating frame of reference [Hill and Marshall(1995); Marshall et al.(1997b); Marshall et al.(1997a)]. The model supports different numerical configurations including hydrostatic and non-hydrostatic simulations. A generic, height like, vertical coordinate allows the fluid equations to be solved in either height or pressure coordinates, making the code amenable to both ocean and atmospheric simulations [Hill et al.(1999); Marshall et al.(2004)]. Adjoint forms of the code are automatically maintained [Heimbach et al.(2002)] supporting advanced assimilation [Stammer et al.(2003)] and dynamical analysis [Marotzke et al.(1999); Hill et al.(2004)]. The configuration used here is hydrostatic and the underlying continuous equations can be written as partial differential equations, in an Eulerian form, that govern the three-dimensional temporal and spatial evolution of velocity  $u$ , potential temperature  $\theta$ , salinity  $s$  and pressure  $g\rho_{ref}\eta + p_{hyd}$  within the fluid

$$\frac{\partial \vec{u}_h}{\partial t} + \nabla_h(g\rho_{ref}\eta + p_{hyd}) = \vec{G}_{\vec{u}_h} \quad (1)$$

$$\frac{\partial \gamma}{\partial t} = G_\gamma \quad \text{where } \gamma = \theta, s, \lambda \quad (2)$$

$$\nabla^2 \eta + \frac{\partial \eta}{\partial t} = \nabla \cdot \widehat{G_{\vec{u}_h}} \quad (3)$$

$$\frac{\partial p_{hyd}}{\partial z} = -g\rho(\theta, s, z) \quad (4)$$

$$\frac{\partial w}{\partial z} = -\nabla_h \cdot \vec{u}_h \quad (5)$$

In (1)  $g\rho_{ref}\eta$  and  $p_{hyd}$  are the surface pressure (due to fluid surface elevation  $\eta$ ) and the internal hydrostatic pressure (due to the mass of the fluid above a particular depth), respectively. The subscript  $_h$  denotes operations in the horizontal plane. Hydrostatic pressure relates to fluid potential temperature,  $\theta$ , salinity,  $s$  and depth  $z$  according to the function  $\rho$  in (4) multiplied by gravity  $g$ . The term  $\vec{\mathbf{G}}_{\vec{\mathbf{u}}_h}$  in (1) incorporates sources and sinks of momentum, due, for example, to surface winds or bottom drag together with terms that represent coriolis forces and the transport of momentum by the fluid flow. Equation (2) captures the time evolution of the ocean heat,  $\theta$ , and salt,  $s$  and tracers,  $\lambda$ , content, respectively. Right-hand-side terms,  $G_{\theta,s,\lambda}$ , include source terms due to surface fluxes and transport by the fluid flow.

To evaluate (1) we require knowledge of the pressure field terms in the fluid,  $g\rho_{ref}\eta + p_{hyd}$ . The hydrostatic part of the pressure field can be found by integrating the fluid density relation  $\rho(\theta, s, z)$  down from the surface, where  $p_{hyd} = 0$ . The surface pressure is the product of the fluid free surface elevation  $\eta$  multiplied by gravity  $g$  and a reference density  $\rho_{ref}$ . Invoking the vertically integrated form of the continuity condition for an incompressible fluid with a free surface,  $\nabla \cdot \hat{\mathbf{u}} + \frac{\partial \eta}{\partial t} = 0$ , yields a diagnostic relation (3) for  $\eta$ . Here  $\hat{\phi}$  indicates a vertical integral from the bottom of the ocean to the surface. The vertical velocity  $w$  is found by integrating (5) vertically from the bottom of the ocean.

Equations (1) and (2) are stepped forward explicitly in time using an Adams-Bashforth procedure that is second order accurate. The equations are discretized in space using a finite volume based technique [Adcroft et al.(1997)] yielding a solution procedure that requires at each time step explicitly evaluated local finite volume computations and an implicit two-dimensional elliptic inversion.

To solve equations (1)-(5) numerically we write them in terms of discrete time levels  $n$ ,  $n + \frac{1}{2}$  and  $n + 1$

$$\frac{\vec{\mathbf{u}}_h^{n+1} - \vec{\mathbf{u}}_h^n}{\Delta t} + \nabla_h(\eta^{n+1} + p_{hyd}^{n+\frac{1}{2}}) = \vec{\mathbf{G}}_{\vec{\mathbf{u}}_h}^{n+\frac{1}{2}} \quad (6)$$

$$\frac{\gamma^{n+1} - \gamma^n}{\Delta t} = G_\gamma^{n+\frac{1}{2}} \quad \text{where } \gamma = \theta, s, \lambda \quad (7)$$

$$\nabla^2 \eta^{n+1} + \frac{\eta^{n+1} - \eta^n}{\Delta t} = \nabla \cdot \widehat{\vec{\mathbf{G}}_{\vec{\mathbf{u}}_h}^{n+\frac{1}{2}}} + \frac{\widehat{\vec{\mathbf{u}}_h^n}}{\Delta t} - p_{hyd}^{n+\frac{1}{2}} \quad (8)$$

$$\frac{\partial p_{hyd}^{n+1}}{\partial z} = -g\rho(\theta^{n+1}, s^{n+1}, z) \quad (9)$$

$$\frac{\partial w^{n+1}}{\partial z} = -\nabla_h \vec{\mathbf{u}}_h^{n+1} \quad (10)$$

The  $n+\frac{1}{2}$  terms in equations (6)-(10) denote centered in time values that are explicitly evaluated from known values at time levels  $n$  and  $n-1$  using Adams-Bashforth extrapolation,  $\phi^{n+\frac{1}{2}} = (\frac{3}{2} + \epsilon)\phi^n - (\frac{1}{2} + \epsilon)\phi^{n-1}$ . Time levels are separated by a dimensional time  $\Delta t$ . The computational algorithm then proceeds as a series of  $n_{final} - n_{initial}$  time steps according to algorithm 1.

Algorithm 1 shapes the computational formulation of MITgcm. Lines 2: and 5: of the algorithm involve integrals along the vertical axis of the discrete domain. Typically these integrals contain dependencies that prevent efficient parallel decom-



```

1: for  $n = n_{initial}$  to  $n_{final}$  do
2:   Finite volume calculations evaluate  $n + \frac{1}{2}$  terms values using Adams-Bashforth
   extrapolation.
3:   Solve a two-dimensional implicit Helmholtz problem that is the finite volume
   discretization of (8) to yield  $\eta^{n+1}$ . This employs a preconditioned conjugate
   gradient solver [Hill and Marshall(1995); Marshall et al.(1998)].
4:   Update horizontal velocity,  $\bar{\mathbf{u}}_h^{n+1}$ , temperature,  $\theta^{n+1}$  and salinity,  $s^{n+1}$ , ac-
   cording to (6) and (7).
5:   Vertically integrate finite volume forms of (9) and (10) to yield  $p_{hyd}^{n+1}$  and
    $w^{n+1}$ .
6: end for

```

**Algorithm 1:** MITgcm dynamical kernel algorithm

position along this axis. Accordingly the parallel formulation takes a global finite volume domain with  $N_x \times N_y \times N_z$  cells in the two horizontal directions,  $x$  and  $y$ , and the vertical direction,  $z$ , respectively, and decomposes it into  $N_{sx} \times N_{sy}$  subdomains each of size  $(S_{nx} + 2 \times O_x) \times (S_{ny} + 2 \times O_y) \times N_z$  such that  $S_{nx} \times N_{sx} = N_x$  and  $S_{ny} \times N_{sy} = N_y$ . The  $O_x$  and  $O_y$  values are overlap region finite volume cells that are added to the boundaries of the subdomains to hold replicated data from neighboring subdomains. The values of  $O_x$  and  $O_y$  are chosen such that the explicit steps 2: and 4: in algorithm 1 can be evaluated locally by each subdomain independent of others. Each computational process integrating forward the MITgcm is then given a static set of one or more subdomains. The process is then responsible for evaluating all the terms required by algorithm 1 for its subdomains and for cooperating with other processes to maintain coherent data in the overlap region finite volume cells. This yields a decomposed computational procedure outlined in algorithm 2. In algorithm 2 a single time-step is split into a series of *compute*, *exchange*, and *sum* phases. *Compute* phases contain only local computation (predominantly arithmetic and associated memory loads and stores) and I/O operations. Performance of *compute* phases is sensitive to the volume of I/O and computation involved, local CPU and memory capabilities of the underlying system, and to the system I/O capacity. *Exchange* phases involve point-to-point communication between neighbor processes. Their performance hinges on the performance of the systems interconnect and inter-process communication software stack. *Sum* phases involve all subdomains collectively combining locally calculated 8-byte floating point values to yield a single global sum. The *sum* phases are sensitive to how system performance for collective communication scales with processor count.

#### 4. RUNNING ON THE ALTIX SYSTEM

The system we are using consists of four interconnected 512 processor SGI Altix machines each with a separate OS image. Each Altix system uses 1.6GHz Itanium 2 processors arranged in pairs and is interconnected with SGI's BX2 NUMALink 4 technology through a front side bus interface. Within each 512-way system NUMALink supports non-uniform access global shared memory [Woodacre et al.(2004)] through memory coherence technology rooted in the distributed shared memory ap-

```

1: for all subdomains of all processes do
2:   Compute terms corresponding to step 2: of algorithm 1.
3: end for
4: Neighbor subdomains exchange overlap regions over  $N_z$  levels.
5: while conjugate gradient solver for equation (8) is not converged do
6:   for all subdomains of all processes do
7:     Compute terms in the preconditioned conjugate gradient algorithm
8:     Neighbor subdomains exchange two-dimensional overlap regions for
       Helmholtz problem.
9:     All processes sum up terms in vector dot products to test for convergence
       and to update the conjugate gradient search directions.
10:   end for
11: end while
12: for all subdomains of all processes do
13:   Compute terms corresponding to steps 4: and 5: of algorithm 1.
14: end for
15: Neighbor subdomains exchange overlap regions over  $N_z$  levels.

```

**Algorithm 2:** MITgcm dynamical kernel decomposed computational procedure for a single time step  $n$ .

proaches of the Stanford DASH project [Lenoski et al.(1992)]. The network fabric within an Altix system is a fat-tree arrangement [Leiserson(1985); Boughton(1997)] based on 8-port router chips. This provides for 64 bisection links on a 512-way system giving a total bisection bandwidth of 400TB/s. Each CPU has memory banks populated with 2GB of 166MHz DDR memory, so that the core memory of a 512-way unit is 1TB. On the Stream benchmark [McCalpin(2005)] a 512-way system achieves a sustained memory bandwidth for the triad value of just over 1TB/s [Ciotti and Nelson(2003)]. Each 512 way system runs a version of the Linux 2.4 kernel with specific SGI modifications that supports scaling to large numbers of shared memory processors [Neuner(2003)].

We scale our runs to 1920 CPUs using a cluster of four 512-way systems interconnected with a single NUMalink fat-tree. The SGI Array Services and Message Passing Toolkit systems are available on the Altix system, providing optimized MPI implementations that can take advantage of the coherent shared memory within each 512-way unit and transparently adapt to non-coherent memory communication, still over NUMalink, between systems. The *exchange* and *sum* phases in the MITgcm algorithm are, therefore, configured to use the systems native MPI implementation for the runs we describe.

The most challenging numerical experiment we have undertaken is a  $\frac{1}{16}^\circ$  near global simulation. In this section we examine the performance of that simulation on the Altix system. The analysis makes use of operation counts based on code inspection as well as runtime processor statistics provided by the pfmon utility [Jarp(2002)] which reads Itanium hardware counters and per MPI process communication statistics available when the `-stats` option of the Message Passing Toolkit is enabled.

The grid spacing for this model configuration is  $\frac{1}{16}^\circ$  longitudinally while latitude

circle spacing is  $\frac{1}{16}^\circ \cos(\text{latitude})$ . The model domain extends from  $78.7^\circ\text{S}$  to  $80^\circ\text{N}$ . This gives a global finite volume grid with grid cells that are  $6.9\text{km} \times 6.9\text{km}$  at the equator and  $1.2\text{km} \times 1.2\text{km}$  at the northern boundary. The number of surface grid cells is just over 25 million, and the configuration has 50 vertical levels, making the total number of cells in all three dimensions just over 1.25 billion. Each of the three dimensional fields that make up the description of the simulation domain and its time evolving state therefore requires 10GB of storage. With the full set of runtime diagnostics that are required for a scientifically useful calculations there are about 180 active three dimensional fields, so that the memory footprint of our running calculation is around 1.8TB. Each time step of algorithm 1 requires computing approximately 3600 floating point operations for each grid point. Aggregated over a total of 1.25 billion grid cells a single time step entails approximately  $4.5 \times 10^{12}$  arithmetic operations.

At each timestep communication is necessary to update overlap regions in the sub-domains on which we compute (see section 3). We use a sub-domain size of  $96 \times 136$  with overlap regions  $O_x = O_y = 3$ . A total of 1920 sub-domains are needed to cover the full domain at this size so that a single three-dimensional *exchange* operation entails transferring a total of 1GB of data. For the configuration used each timestep involves eighteen such operations, corresponding to steps 4: and 15: in algorithm 2.

Additionally an average of forty iterations in the conjugate gradient solver are required at each timestep. Each iteration entails a pair of two-dimensional exchange operations (each involving 20MB of data transfer) and at least one global sum of an 8-byte floating point value from each sub-domain.

A large number of diagnostics are saved to allow analysis. In addition to basic state information high-frequency outputs of two-dimensional sea-surface elevation and bottom pressure that can be related to satellite measurements are saved along with three-dimensional product terms that cannot be derived from their separate components. The list of possible outputs is application dependent but at a minimum the I/O requirement is 300GB per month of integration.

The 1.8TB memory footprint of our calculation is larger than the main memory of a single 512-way box so for this study we examine processor counts of 960, 1440 and 1920 which correspond to running across two, three and four 512-way Altix boxes. We make use of a dedicated sets of CPUs and employ the Message Passing Toolkit `MPI_DSM_DISTRIBUTE` flags to ensure that when an MPI process is assigned to an Altix node, the CPU and memory on that node bind to that process.

Figure 2 shows the scaling behavior of two key communication primitives used in MITgcm, the *sum* operation and the two-dimensional form of the *exchange* operation. These times were obtained from an isolated kernel benchmark, taken from the MITgcm code, because the full,  $\frac{1}{16}^\circ$ , configuration will not fit on smaller numbers of processors. The times, shown in microseconds, increase with processor count for both the *sum* operation and the two-dimensional *exchange*. However, it is encouraging to note that the increase in time when going between different numbers of Altix boxes is comparable to the increase associated with going from 1 to 512 CPUs. The *exchange* times stay relatively flat because although more data is exchanged between processors as the number of CPUs increases the available network




Fig. 2. Performance of key primitives used on the  $\frac{1}{16}^\circ$  resolution simulation. The exchange times are for a sub-domain of size  $96 \times 136$  with  $O_x = O_y = 1$

bandwidth increases in proportion with CPU count.

Figure 3 shows the overall scaling and performance of a full simulation. Two scenarios are shown one with a modest I/O load (the upper line) and one with a very high I/O load (the lower line). The communication costs in figure 2 are only a small fraction of the total time in both cases. For the modest I/O case the peak performance at 1920 processors is sufficient to achieve 45 days of simulation per wall-clock day with an aggregate performance of 1.39 TFlop/s. The modest I/O case also scales well from 960 processors. For the intensive I/O scenario performance drops so that maximum throughput on 1920 processors is around 17 simulated days per wall-clock day with an aggregate performance of around 460 GFlop/s. For the intensive I/O scenario the scaling from 960 processors falls below the linear relation (the dotted line). However, at present, the MITgcm I/O routines being used have not been optimized for the Altix systems and so we expect the performance of the intensive I/O scenario to improve significantly. Overall these results are encouraging and are sufficient to enable us to proceed with some preliminary numerical experiments.

## 5. PRELIMINARY RESULTS

Using the Altix configuration described earlier, we undertook a series of numerical simulations at  $\frac{1}{4}^\circ$ ,  $\frac{1}{8}^\circ$ , and  $\frac{1}{16}^\circ$  resolutions. The three cases are all configured as described in section 4 and, apart from differing resolutions, the model configurations are the same. Figures 4, 5 and 6 show significant changes in solution with resolution. Each plot shows the change in simulated sea-surface height over the same month.




Fig. 3. Overall scaling and performance of the  $\frac{1}{16}^\circ$  resolution simulation on 960, 1440 and 1920 processors.

The plots capture changes due to eddy activity over a single month. Changes with resolution occur in the regions of the oceans where eddies are prevalent (such as the Gulf Stream, the Kurishio, the Aghullas, the Drake Passage and the Antarctic Circumpolar Current). For example in the Gulf Stream region the upper panel of figure 4 ( $\frac{1}{4}^\circ$ ) shows a relatively small area of vigorous sea-surface height changes. The middle and lower panels ( $\frac{1}{8}^\circ$  and  $\frac{1}{16}^\circ$  respectively) show more extensive areas of changes but there are big shifts between the  $\frac{1}{8}^\circ$  and the  $\frac{1}{16}^\circ$  results.

Figure 5 and 6 are closeups of the Gulf Stream and Drake Passage. Key behaviors like how tightly waters “stick” to the coast, or how far energetic eddies penetrate the ocean interior change significantly between resolutions.

## 6. CONCLUSIONS

At first glance the three different resolution runs show significant differences. There does, however, seem to be smaller change between the  $\frac{1}{8}^\circ$  and  $\frac{1}{16}^\circ$  simulations. A next step is to undertake a fourth series of runs at even higher resolution,  $\frac{1}{20}^\circ$  or  $\frac{1}{32}^\circ$ . Formally quantifying the changes between these runs would provide important information on whether ocean models are reaching numerically converged solutions.

Performance on the Altix shows it is well suited for addressing these questions. We monitored our code to be achieving about 722 Mflop/s/cpu when running on 1920 processors. This is 14% of the per CPU Top500 number achieved on the system [Meuer and Dongarra(2005)]. Our code consists of predominantly BLAS1 class operations and cannot exploit the level of cache reuse that the Top500 Linpack benchmark achieves. The scaling we found across multiple Altix systems is

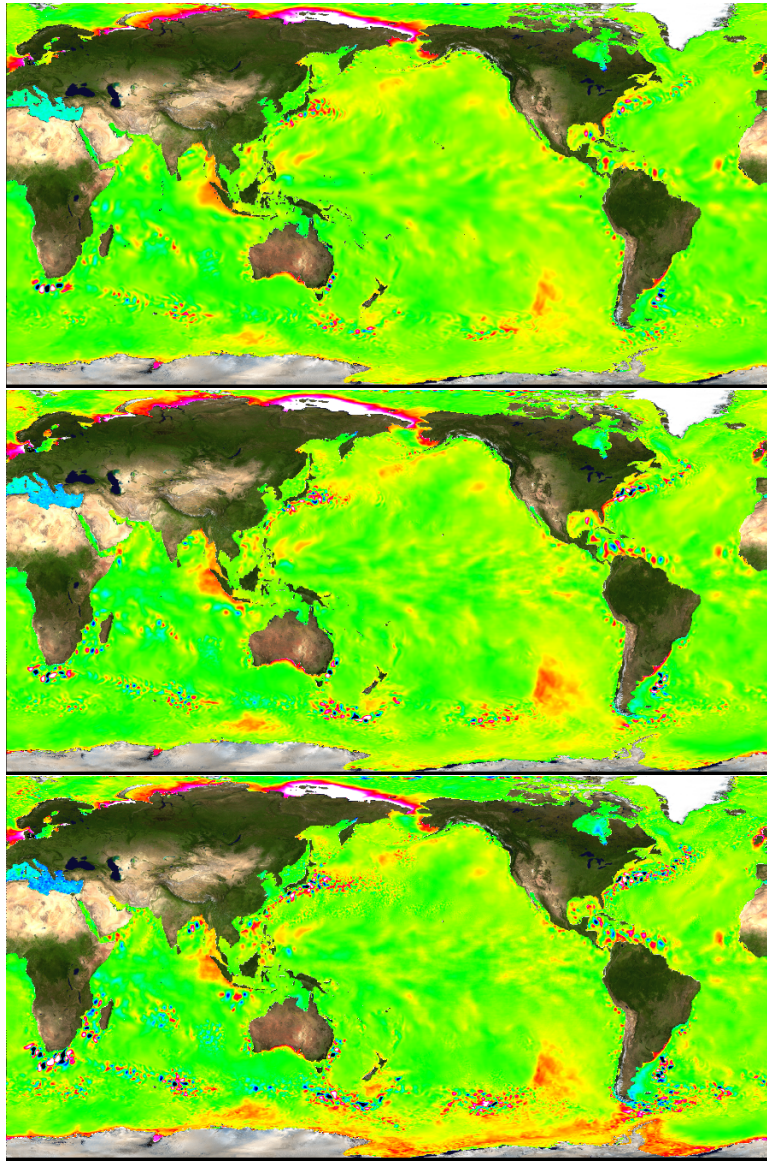


Fig. 4. One month sea-surface height differences. Upper  $\frac{1}{4}^\circ$ , middle  $\frac{1}{8}^\circ$ , lower  $\frac{1}{16}^\circ$ . Color scale -0.5m to 0.5m.

encouraging and suggests that configurations that span eight or more Altix boxes, and that would therefore support  $\frac{1}{20}^\circ$  and higher resolution, are today within reach.

#### REFERENCES

- A. Adcroft, C. Hill, and J. Marshall. Representation of topography by shaved cells in a height coordinate ocean model. pages, 2293-2315, 1997.
- ACM Transactions on Computational Logic, Vol. V, No. N, April 2005.






Fig. 5. Gulf Stream region sea-surface height difference plots at different resolutions for the same time period as figure 4. Left panel  $\frac{1}{4}^\circ$ , middle panel  $\frac{1}{8}^\circ$ , right panel  $\frac{1}{16}^\circ$ . Color scale -0.125m to 0.125m.




Fig. 6. Drake Passage region sea-surface height difference plots at different resolutions for the same time period as figure 4. Left panel  $\frac{1}{4}^\circ$ , middle panel  $\frac{1}{8}^\circ$ , right panel  $\frac{1}{16}^\circ$ . Color scale -0.125m to 0.125m.

- C.W. Boening and R. Budich. Eddy dynamics in a primitive equation model: sensitivity to horizontal resolution and friction. *J. Phys. Oceanogr.*, 22:361–381, 1992.
- G.A. Boughton. Arctic Switch Fabric. In *Proceedings of the Second International Workshop on Parallel Computer Routing and Communication*, pages, 65-74. Springer-Verlag, Berlin (Germany), 1997.
- K. Bryan and W. Holland. A high resolution simulation of the wind and thermohaline driven circulation of the North Atlantic Ocean. pages, 99-115. Hawaii Institute of Geophysics Manoa, 1989.
- D. Chelton, R. DeSzoeke, and M. Schallax. Geophysical variability of the first baroclinic rossby radius of deformation. *J. Phys. Oceanogr.*, 28:433–460, 1998.
- B. Ciotti and B. Nelson. STREAM benchmark of 512 CPU Altix. <http://www.cs.virginia.edu/stream/>, 2003.
- T. Haine and J. Marshall. Gravitational, symmetric and baroclinic instability of the ocean mixed layer. *J. Phys. Oceanogr.*, 29:634–658, 1998.
- P. Heimbach, C. Hill, and R. Giering. Automatic generation of efficient adjoint code for a parallel navier-stokes solver. In J.J. Dongarra, P.M.A. Soot and C.J.K. Tan, editor, *Computational Science – ICCS 2002*, pages, 1019-1028. Springer-Verlag, Berlin (Germany), 2002.
- C. Hill and J. Marshall. Application of a Parallel Navier-Stokes Model to Ocean Circulation. pages, 545–552, 1995.
- C. Hill, A. Adcroft, D. Jamous, and J. Marshall. A strategy for terascale climate modeling. In *ACM Transactions on Computational Logic*, Vol. V, No. N, April 2005.

- Proceedings of the Eight ECMWF Workshop on the Use of Parallel Processors in Meteorology*, 1999.
- C. Hill, V. Bugnion, M. Follows, and J. Marshall. Evaluating carbon sequestration efficiency in an ocean circulation model by adjoint sensitivity analysis. *J. Geophys. Res.*, 109(C11005), 2004. doi:10.1029/2002JC001598.
- H.E. Hurlburt and P.J. Hogan. Impact of 1/8 to 1/64 resolution on gulf stream model-data comparisons in basin-scale subtropical atlantic ocean models. *Dyn. Atmos and Oceans*, 32:283–329, 2000.
- S. Jarp. A Methodology for using the Itanium 2 Performance Counters for Bottleneck Analysis. HP Labs, 2002.
- Helen Jones and John Marshall. Restratification after deep convection. *J. Phys. Oceanogr.*, 27:2276–2287, 1997.
- C.E. Leiserson. Fat-Trees: Universal Networks for Hardware-Efficient Supercomputing. *IEEE Transactions on Computers*, C34:892–901, October 1985.
- D. Lenoski, J. Laudon, K. Gharachorloo, W. Weber, A. Gupta, J. Hennessy, M. Horowitz, , and M. Lam. The Stanford Dash multiprocessor. pages, 63–79, March 1992.
- M.E. Maltrud and J.L. McClean. An eddy resolving global 1/10 ocean simulation. *Ocean Modeling*, 8:31–54, 2005.
- J. Marotzke, R. Giering, K. Q. Zhang, D. Stammer, C. Hill, and T. Lee. Construction of the Adjoint MIT Ocean General Circulation Model and Application to Atlantic Heat Transport Sensitivity. *J. Geophys. Res.*, 104(C12):29,529–29,547, 1999.
- J. Marshall, A. Adcroft, C. Hill, L. Perelman, and C. Heisey. A finite-volume, incompressible navier stokes model for studies of the ocean on parallel computers. *J. Geophys. Res.*, 102, C3:5,753–5,766, 1997a.
- J. Marshall, C. Hill, L. Perelman, and A. Adcroft. Hydrostatic, quasi-hydrostatic and nonhydrostatic ocean modeling. *J. Geophys. Res.*, 102, C3:5,733–5,752, 1997b.
- J. Marshall, H. Jones, and C. Hill. Efficient ocean modeling using non-hydrostatic algorithms. *Journal of Marine Systems*, 18:115–134, 1998.
- J. Marshall, A. Adcroft, J-M. Campin, C. Hill, and A. White. Atmosphere-ocean modeling exploiting fluid isomorphisms. *Monthly Weather Review*, 132(12):2882–2894, 2004.
- Y. Masumoto et al. A fifty-year eddy-resolving simulaiton of the world ocean. *Journal of the Earth Simulator*, 1:35–56, April 2004.
- J. McCalpin. The STREAM benchmark web site. <http://www.cs.virginia.edu/stream/>, 2005.
- J. McClean, P.M. Poulain, J.W. Pelton, and M.E. Maltrud. Eulerian and Lagrangian statistics from surface drifters and a high-resolution POP simulation in the North Atlantic. *J. Phys. Oceanogr.*, 32:2473–2491, 2002.
- H. Meuer and J. Dongarra. The Top 500 benchmark web site. <http://www.top500.org>, 2005.
- S. Neuner. Scaling Linux to New Heights: the SGI Altix 3000 System. *Linux Journal*, 2003.
- A.J.G Nurser and J.W.Zhang. Eddy-induced mixed layer shallowing and mixed layer/thermocline exchange. *J. Geophys. Res.*, 105:21851–21868, 2000.
- A. Oschillies. Improved representation of upper-ocean dynamics and mixed layer depths in a model of the North Atlantic on switching from eddy-permitting to eddy-resolving grid resolution. *J. Phys. Oceanogr.*, 32:2277–2298, August 2002.
- A. Semtner and R. Chervin. Ocean general circulation from a global eddy-resolving model. *J. Geophys. Res.*, 97:5493–5550, 1992.
- R.D. Smith, M.E. Maltrud, F. Bryan, and M.W. Hecht. Numerical simulation of the North Atlantic at 1/10. *J. Phys. Oceanogr.*, 30:1532–1561, 2000.
- D. Stammer, C. Wunsch, R. Giering, C. Eckert, P. Heimbach, J. Marotzke, A. Adcroft, C. Hill, , and J. Marshall. Volume, heat, and freshwater transports of the global ocean circulation 1993–2000, estimated from a general circulation model constrained by World Ocean Circulation Experiment (WOCE) data. *J. Geophys. Res.*, 108(C1):3007–3029, 2003.
- M. Woodacre, D. Robb, D. Roe, and K. Feind. The SGI Altix 3000 Global Shared-Memory Architecture. White paper, SGI, Mountain View(CA), USA, 2004.